arXiv:1611.01393v1 [cs.SI] 4 Nov 2016

Hierarchical Overlapping Clustering of Network Data Using Cut Metrics

Fernando Gama, Santiago Segarra, and Alejandro Ribeiro

***Abstract*—A novel method to obtain hierarchical and overlapping clusters from network data – i.e., a set of nodes endowed with pairwise dissimilarities – is presented. The introduced method is *hierarchical* in the sense that it outputs a nested collection of groupings of the node set depending on the resolution or degree of similarity desired, and it is *overlapping* since it allows nodes to belong to more than one group. Our construction is rooted on the facts that a hierarchical (non-overlapping) clustering of a network can be equivalently represented by a finite ultrametric space and that a convex combination of ultrametrics results in a cut metric. By applying a hierarchical (non-overlapping) clustering method to multiple dithered versions of a given network and then convexly combining the resulting ultrametrics, we obtain a cut metric associated to the network of interest. We then show how to extract a hierarchical overlapping clustering structure from the aforementioned cut metric. Furthermore, the so-called overlapping function is presented as a tool for gaining insights about the data by identifying meaningful resolutions of the obtained hierarchical structure. Additionally, we explore hierarchical overlapping quasi-clustering methods that preserve the asymmetry of the data contained in directed networks. Finally, the presented method is illustrated via synthetic and real-world classification problems including handwritten digit classification and authorship attribution of famous plays.**

***Index Terms*—Clustering, Network theory, Cut metrics, Hierarchical clustering, Covering, Dithering.**

I. Introduction

Consider a dataset where each element can be represented as a node in a network with a pairwise dissimilarity function. In this setting, the general objective of clustering is to group those nodes that are more similar to each other than to the rest, according to the relationship established by the dissimilarity function [1], [2]. Clustering and its generalizations are ubiquitous tools since they are used in a wide variety of fields such as psychology [3], social network analysis [4], political science [5], neuroscience [6], among many others [7], [8].

Traditional clustering methods provide only one partitioning of the node set in such a way that each data point belongs to one and only one block of the partition. An important limitation of these traditional clustering methods is that the dataset may present a complex data structure at several resolutions or levels of similarity, and outputting only one partition may not be adequate in portraying the different grouping degrees that may be present. *Hierarchical* clustering solves this issue by providing a collection of partitions, indexed by a resolution parameter, that can be set by the user to determine the extent up to which nodes are considered similar or different [9]. In other words, each partition in the collection reflects a different degree of similarity between the nodes, ranging from considering all nodes different, to considering all nodes similar. Examples of hierarchical clustering methods are UPGMA [10], Ward's method [11], complete [12] or single linkage [13].

However, hierarchical methods still have a major limitation, namely, that nodes are assigned to one and only one category or cluster. There are many situations in which this assignment does not constitute a reasonable approach. Specifically, some nodes might inherently have traits of more than one group and hence have similarities to multiple clusters of nodes that otherwise would be considered dissimilar [14], [15]. Two classes of methods have been proposed to overcome this issue. The first class corresponds to the so-called soft or fuzzy clustering methods that allow the allocation of a node to multiple clusters by assigning to each node a membership degree or probability of belonging to every cluster [16]. Nevertheless, these methods subscribe to the idea that nodes have more affinity to either one or another subset, as illustrated by their membership degree; or that they belong to only one group but the setting of the problem is not rich enough to determine to which one, thus assigning a probability of belonging to different subsets. Fundamentally, these methods do not contemplate the idea that a node can intrinsically be part in equal terms of more than one cluster. The second class of methods that attempt to solve this issue is the class of *overlapping* clustering methods, which perform a non-hierarchical deterministic assignment of nodes to more than one subset [17]. While overlapping methods enable nodes to belong to more than one cluster, they are still myopic – just like traditional clustering methods – to multiple affinity resolutions.

In this paper, a *hierarchical overlapping* clustering method is proposed. Our method outputs a collection of groupings where each of them corresponds to different resolutions of similarity between the nodes (hierarchical), while allowing nodes in each grouping to deterministically belong to more than one cluster (overlapping). Essentially, the method is derived by generalizing the concepts of ultrametrics, equivalence relations, and dendrograms – typical of hierarchical (non-overlapping) clustering methods – to those of cut metrics [18], [19], tolerance relations [20], and nested collections of coverings. More precisely, we generate

Work in this paper is supported by NSF CCF 1217963. F. Gama and A. Ribeiro are with the Department of Electrical and Systems Engineering, University of Pennsylvania. S. Segarra is with the Institute for Data, Systems, and Society, Massachusetts Institute of Technology. Emails: fgama@seas.upenn.edu, segarra@mit.edu, aribeiro@seas.upenn.edu.

a family of networks derived from dithering multiple times the dissimilarity function of a given network. We then proceed to apply a hierarchical (non-overlapping) clustering method to each of the dithered networks, resulting in a collection of ultrametrics. From this collection we obtain a cut metric related to the network of interest by means of a convex combination of the ultrametrics. Finally, based on the cut metric we construct a tolerance relation that connects the nodes that are at a distance lower than a pre-specified resolution, which we then use to derive a covering of the node set. By considering all possible resolution levels, we obtain a nested collection of coverings.

In Section II, the concepts of partition, equivalence relation, and dendrogram are defined, and their relations to both traditional and hierarchical (non-overlapping) clustering methods are established. These concepts are then respectively generalized to those of covering, tolerance relation, and nested collection of coverings. This generalization is achieved by dropping the requirements that prevent a node from belonging to more than one group. Moreover, a formal definition of a hierarchical overlapping clustering method is presented. Section III relates the introduced concepts with concrete finite metric constructions. More specifically, ultrametrics are defined and directly linked to hierarchical clustering methods (Theorem 1). We then introduce cut metrics and explain how they can be used to obtain a nested collection of coverings (Theorem 2). The proposed algorithm to obtain cut metrics from a given network (Algorithm 1), and hence nested collection of coverings, is detailed in Section IV. The algorithm is grounded on the fact that a convex combination of ultrametrics yields a cut metric (Proposition 1) and the concept of dithering [21], that can be used to obtain multiple noisy versions of the given network. Also, quasi-clustering methods that allow for a hierarchical structure and overlapping nodes are explored in Section V. These methods are particularly useful when there is a need to preserve in the grouping structure the asymmetric nature of directed networks. Section VI illustrates the implementation of the proposed method in synthetic experiments to show that the hierarchical overlapping clustering algorithm yields sensitive results that correspond to intuition. Finally, in Section VII, the proposed method is applied to the unsupervised classification problem of determining handwritten digits [22] and the authorship attribution problem of identifying plays that have been co-authored [23]. Concluding remarks in Section VIII close the paper.

II. Preliminaries

Let $N = (X, A_X)$ be a network defined by a finite set of nodes X and a dissimilarity function $A_X : X \times X \to \mathbb{R}_+$ that measures how different two nodes are. This function satisfies that, for any two nodes $x, x' \in X$, $A_X(x, x') \geq 0$ with $A_X(x, x') = 0$ if and only if $x = x'$. A_X is a dissimilarity function in the sense that the greater the value of A_X, the more different two nodes are. Observe that the dissimilarity function need not satisfy the triangle inequality nor be symmetric, that is, we may have that $A_X(x, x') \neq A_X(x', x)$. This is particularly useful when modeling uneven levels of influence among nodes. We denote by $\mathcal{N}$ the set of all possible networks.

Definition 1. *A partition $P_X = \{B_1, \ldots, B_m\}$ is a collection of subsets of X such that $\cup_{i=1}^m B_i = X$ and $B_i \cap B_j = \emptyset$ for $i \neq j$, $i, j = 1, \ldots, m$. Denote by $\mathcal{P}$ the set of all possible partitions.*

Definition 2. *An equivalence relation $\sim$ is a binary set relation that, for any $x, x', x'' \in X$, satisfies the properties of reflexivity ($x \sim x$), symmetry ($x \sim x'$ if and only if $x' \sim x$), and transitivity (if $x \sim x'$ and $x' \sim x''$, then $x \sim x''$).*

Observe that partitions (Def. 1) uniquely define equivalence relations (Def. 2) as follows: $x \sim x'$ if and only if $x, x' \in B_i$ for some $B_i \in P_X$, $i = 1, \ldots, m$. Equivalently, equivalence relations uniquely define partitions [24]. An equivalence relation obtained from a specific partition P_X will be denoted by $\sim_{P_X}$.

With these definitions in place, a (traditional) clustering method $\mathcal{G}$ can be defined as a structure-preserving map from the set of networks to the set of partitions, $\mathcal{G} : \mathcal{N} \to \mathcal{P}$. This map is structure preserving in the sense that the output partition is defined over the same node set X specific to the input network. That is, $\mathcal{G}(N) = P_X = \{B_1, \ldots, B_m\}$ for $N = (X, A_X)$ with $\cup_{i=1}^m B_i = X$. Intuitively, a desirable clustering method $\mathcal{G}$ is one in which the nodes contained in each subset B_i are determined by the dissimilarity function A_X so that similar nodes are grouped together.

In many situations, having only one partition as the output of a clustering method may not be appropriate. It might be desirable to output several partitions that depend on a resolution parameter that specifies the affinity required between two nodes to be deemed as similar.

Definition 3. *A dendrogram $D_X = \{D_X(\delta), \delta \geq 0\}$ is a collection of partitions, where δ is a resolution parameter and $D_X(\delta)$ is a partition of the node set X, satisfying the following two properties:*

(i) $D_X(0) = \{\{x\}, x \in X\}$, and there exists a $\delta_{\max}$ such that for all $\delta \geq \delta_{\max}$, $D_X(\delta) = \{X\}$;

(ii) it is a nested collection of partitions, that is, for $\delta \leq \delta'$, if $x \sim_\delta x'$, then $x \sim_{\delta'} x'$ for all $x, x' \in X$;

together with a technical condition for right-continuity [25].

Note that, in an effort to ease the exposition, $x \sim_\delta x'$ is being used as a shorthand notation for $x \sim_{D_X(\delta)} x'$. Observe that property (i) encodes the fact that at resolution $\delta = 0$ all nodes are considered to be different from each other, thus each node forms a singleton cluster, whereas for sufficiently large δ all nodes are considered to be similar and, hence, clustered in a single block. Property (ii) enforces the agglomeration to occur in a nested fashion in the sense that if two nodes belong to the same cluster at a specific resolution then they will remain co-clustered for any larger resolution. Denoting by $\mathcal{D}$ the set of all possible dendrograms, we define a hierarchical clustering method $\mathcal{H}$ as follows.

Figure 1. Dumbbell network with dissimilarities given by the Euclidean distance between points.

Definition 4. *A hierarchical clustering method $\mathcal{H}$ is a structure-preserving map from the set of networks to the set of dendrograms*

$$\mathcal{H} : \mathcal{N} \to \mathcal{D}. \tag{1}$$

A limitation of both traditional $\mathcal{G}$ and hierarchical $\mathcal{H}$ clustering methods is that each node is required to belong to one and only one cluster at any given resolution. This requirement stems from the nonintersecting nature of the blocks that form a partition. In many problems, however, some nodes may inherently have traits of more than group, making it reasonable for them to belong to more than one cluster. For instance, consider the dumbbell network in Fig. 1 where the dissimilarity between two nodes is given by their Euclidean distance. It is intuitive that the point clouds on each side should form separate clusters. However, it is unclear if, e.g., *all* of the handle (bridge) should be a separate cluster as it is not unreasonable to assign its borders to the respective point clouds.

Throughout the paper, we develop a clustering method that enables both deterministic overlapping of nodes and hierarchical outcomes. This is achieved by generalizing the concepts of equivalence relations, partitions, and dendrograms to those of tolerance relations, coverings, and nested collections of coverings, respectively; see Table I.

A. Generalizing concepts: Coverings and tolerance relations

The first step in order to allow a node to belong to more than one cluster is to drop the nonintersecting requirement of partitions.

Definition 5. *A covering $Q_X = \{C_1, \dots, C_m\}$ of the node set X is a collection of subsets such that $\cup_{i=1}^m C_i = X$, but not necessarily $C_i \cap C_j = \emptyset$ for $i \neq j$. Denote by $\mathcal{Q}$ the set of all possible coverings.*

Note that partitions (Def. 1) are particular cases of coverings (Def. 5), i.e., $\mathcal{P} \subset \mathcal{Q}$. Furthermore, a non-hierarchical overlapping clustering method $\mathcal{M}$ is defined as a structure-preserving map from the set of networks to the set of coverings $\mathcal{M} : \mathcal{N} \to \mathcal{Q}$. Related to the concept of covering, we introduce the notion of a tolerance relation.

Definition 6. *A tolerance relation $\leftrightarrow$ is a binary set relation that, for any $x, x' \in X$, satisfies the properties of reflexivity ($x \leftrightarrow x$), and symmetry ($x \leftrightarrow x'$ if and only if $x' \leftrightarrow x$), but is not necessarily transitive.*

Observe that a tolerance relation is a generalization of an equivalence relation (Def. 2) obtained by dropping transitivity. A tolerance relation defined on a node set X can be represented by an unweighted and undirected graph where an edge exists between nodes x and x' if and only if $x \leftrightarrow x'$. There are two established ways of inducing a covering from a tolerance relation based on this graph [20]. The first methodology is referred to as covering by classes, wherein each cover consists of the neighbor set of each node, together with the node itself, and where all covers that are completely contained within a larger one are removed. The other established procedure is denominated covering by blocks, where each block of the covering is given by the maximal cliques of the graph. For the remainder of the paper, we adopt this latter procedure as the canonical way of inducing coverings from tolerance relations. Note that, unlike the case of equivalence relations and partitions, while any tolerance relation can induce a covering (by blocks) not all coverings of a node set X can be induced by a tolerance relation [20, Theorem 2]. Leveraging the connection between tolerance relations and coverings, we extend the notion of a dendrogram into the realm of overlapping clusters.

Table I: Parallelism between Hierarchical clustering $\mathcal{H}$ and Hierarchical Overlapping clustering $\mathcal{O}$.

Method	Hierarchical $\mathcal{H}$ (Def. 4)	Overlapping $\mathcal{O}$ (Def. 8)
Metric	Ultrametric $u_X(x, x')$ (Def. 9)	Cut Metric $c_X(x, x')$ (Def. 10)
Relation	Equivalence $\sim$ (Def. 2)	Tolerance $\leftrightarrow$ (Def. 6)
Grouping	Partition $P_X = \{B_i\}$ (Def. 1)	Covering $Q_X = \{C_i\}$ (Def. 5)
Hierarchy	Dendrogram D_X (Def. 3)	Nested Covering K_X (Def. 7)
Construction	Theorem 1	Theorem 2

Definition 7. *A nested covering $K_X = \{K_X(\delta), \delta \geq 0\}$ is a collection of coverings, where δ is a resolution parameter and $K_X(\delta)$ is a covering induced by a tolerance relation $\leftrightarrow_\delta$, satisfying the following properties:*

(i) $K_X(0) = \{\{x\}, x \in X\}$ and there exists a $\delta_{\max}$ such that for all $\delta \geq \delta_{\max}$, $K_X(\delta) = \{X\}$;

(ii) For $\delta \leq \delta'$, if $x \leftrightarrow_\delta x'$, then $x \leftrightarrow_{\delta'} x'$ for all $x, x' \in X$.

Notice the resemblance between properties (i) and (ii) above and those in Def. 3. In this way, we extend the notion of nestedness from its intuitive definition for (non-overlapping) clusters to the setting of overlapping clusters. Denoting by $\mathcal{K}$ the set of all nested coverings, we define hierarchical overlapping clustering methods as follows.

Definition 8. *A hierarchical overlapping clustering method $\mathcal{O}$ is a structure-preserving map from the set of networks to the set of nested coverings*

$$\mathcal{O} : \mathcal{N} \to \mathcal{K}. \tag{2}$$

Given a network $N = (X, A_X)$, a clustering method $\mathcal{O}$ outputs a structure $\mathcal{O}(N) = K_X$ which takes into account different resolution levels of similarity and, at each resolution, allows overlapping clusters. Building on metric concepts developed in Section III, we present the practical construction of one such method $\mathcal{O}$ in Section IV.

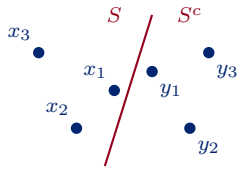

Figure 2. A cut semimetric on the node set $X = \{x_1, x_2, x_3, y_1, y_2, y_3\}$. A cut that partitions the space into S and S^c is drawn. The resulting cut semimetric is $\delta_S(x_i, x_j) = \delta_S(y_i, y_j) = 0$ and $\delta_S(x_i, y_j) = 1$ for all $i, j = 1, 2, 3$.

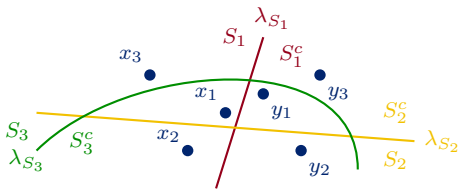

Figure 3. Construction of a cut metric c_X. Consider the node set $X = \{x_1, x_2, x_3, y_1, y_2, y_3\}$. The first cut separates $\{x_i\}$ from $\{y_i\}$, $i = 1, 2, 3$ and is assigned a dissimilarity λ_{S_1}. The cuts associated with S_i for $i = 2, 3$ separate $\{x_i, y_i\}$ from the rest of X and are assigned dissimilarity λ_{S_i}. Finally, the cut metric is obtained as $c_X = \lambda_{S_1}\delta_{S_1} + \lambda_{S_2}\delta_{S_2} + \lambda_{S_3}\delta_{S_3}$.

III. Nested coverings derived from cut metrics

Hierarchical clustering methods output dendrograms, which are nested collections of partitions. However, dendrograms are cumbersome to handle mathematically so their equivalence to ultrametric spaces proves to be useful [13].

Definition 9. *An ultrametric is a function $u_X : X \times X \to \mathbb{R}_+$ that, for any $x, x', x'' \in X$ satisfies*

(i) $u_X(x, x') \geq 0$ *and* $u_X(x, x') = 0$ *if and only if* $x = x'$*;*
(ii) $u_X(x, x') = u_X(x', x)$*;*
(iii) $u_X(x, x'') \leq \max\{u_X(x, x'), u_X(x', x'')\}$.

Ultrametrics can be interpreted as metrics that satisfy a more stringent version of the triangle inequality [Property (iii)] where the addition in the regular triangle inequality is replaced by a maximization. Dendrograms can be equivalently represented as ultrametrics, as shown in the following theorem.

Theorem 1. *Given an ultrametric $u_X(x, x')$ defined over the node set X, define a collection of equivalence relations as follows*

$$u_X(x, x') \leq \delta \iff x \sim_\delta x'. \tag{3}$$

The resulting collection of equivalence relations is such that it generates a dendrogram D_X.

Likewise, given a dendrogram D_X that uniquely determines a collection of equivalence relations $\sim_\delta$ defined on the node set X, an ultrametric u_X can be obtained as

$$u_X(x, x') := \min\{\delta : x \sim_\delta x'\}. \tag{4}$$

Proof: See [25, Theorem 1]. ■

Theorem 1 establishes a structure-preserving bijection between ultrametrics and dendrograms. Consequently, we may reinterpret hierarchical clustering methods as structure-preserving maps from the set of networks to the set $\mathcal{U}$ of all possible ultrametrics, i.e., $\mathcal{H} : \mathcal{N} \to \mathcal{U}$ (cf. Def. 4).

With the objective of deriving a relation between nested coverings K_X and a metric structure akin to that between dendrograms and ultrametrics, the notion of cut metrics is introduced. Let S be a subset of a given node set X and define a cut semimetric $\delta_S(x, x')$ as

$$\delta_S(x, x') = \mathbb{I}\{S \cap \{x, x'\} \neq \emptyset\}\mathbb{I}\{S^c \cap \{x, x'\} \neq \emptyset\} \tag{5}$$

where $\mathbb{I}\{\}$ is the indicator function. Essentially, the cut semimetric partitions the node set into two blocks S and S^c and assigns a dissimilarity of 1 to nodes that are in different sides of this partition, and 0 to nodes that fall in the same side; see Fig. 2.

This cut semimetric can also be understood as being induced by a binary classifier which partitions the node set in two and assigns a unit dissimilarity to nodes belonging to different categories. If multiple cuts, or partitions, are considered, then a cut metric c_X can be obtained as the combination of all the associated dissimilarities; see Fig. 3.

Definition 10. *A cut metric $c_X : X \times X \to \mathbb{R}_+$ is a metric that can be written as a nonnegative linear combination of cut semimetrics*

$$c_X(x, x') = \sum_{S \subseteq X} \lambda_S \delta_S(x, x') \tag{6}$$

with $\lambda_S \geq 0$ and where the sum ranges over all possible subsets of X.

Note that $\lambda_{S'} = 0$ for some S' amounts to the associated cut not being considered in the construction of c_X. It follows from Def. 10 that c_X is a metric, thus, in particular it satisfies that for any $x, x' \in X$, $c_X(x, x') = c_X(x', x) \geq 0$ (symmetry and nonnegativity) and $c_X(x, x') = 0$ if and only if $x = x'$ (identity). We denote by $\mathcal{V}$ the set of cut metrics.

In the same way that ultrametrics induce dendrograms (Theorem 1), cut metrics induce nested collections of coverings, as shown next.

Theorem 2. *Let X be a set of nodes and let c_X be a cut metric defined on X. Let $K_X = \{K_X(\delta), \delta \geq 0\}$ be a collection of coverings. If, for each $\delta \geq 0$, the corresponding covering $K_X(\delta)$ is obtained from the tolerance relation given by*

$$c_X(x, x') \leq \delta \Longrightarrow x \leftrightarrow_\delta x' \tag{7}$$

then K_X is a nested collection of coverings as defined in Def. 7.

Proof: Observe that (7) defines a valid tolerance relation since $c_X(x, x) = 0$ and $c_X(x, x') = c_X(x', x)$ imply reflexivity and symmetry of $\leftrightarrow_\delta$, respectively. Using the relation $\leftrightarrow_\delta$ in (7) to define a collection of coverings by blocks $K_X(\delta)$ [20], it can be shown that K_X satisfies properties (i) and (ii) in Def. 7. To be more specific, $K_X(0) = \{\{x\}, x \in X\}$ follows from the fact that $c_X(x, x') = 0$ if and only if $x = x'$ and, since there exists a

$\delta_{\max}$ such that $c_X(x,x') \leq \delta_{\max}$ for all $x, x' \in X$, it follows that $K_X(\delta) = \{X\}$ for all $\delta \geq \delta_{\max}$. Finally, observe that K_X is nested since, for any $\delta < \delta'$, $c_X(x,x') \leq \delta < \delta'$ implies that $x \leftrightarrow_\delta x' \Rightarrow x \leftrightarrow_{\delta'} x'$. ■

Notice that (7) is similar to the connection between ultrametrics and equivalence relations in (3) but in a unidirectional fashion; see Remark 1. For $0 < \delta < \delta_{\max}$ the resulting tolerance relations define coverings in such a way that nodes might belong to more than one cluster. If there exists at least one such node, it is said that there is *overlap*. Consequently, we are interested in computing the number of overlapping nodes for each value of δ.

Definition 11. *Let $K_X = \{K_X(\delta), \delta \geq 0\}$ be a nested collection of coverings defined over node set X and where each covering $K_X(\delta)$ is comprised of $m(\delta)$ blocks or covers, i.e. $K_X(\delta) = \{C_1, \ldots, C_{m(\delta)}\}$. Then, the overlapping function $f_{ol} : \mathbb{R}_+ \to \mathbb{Z}_+$ is given by*

$$f_{ol}(\delta) = \sum_{k=1}^{n} \mathbb{I}\left\{x_k \in C_i \cap C_j,\ i \neq j,\ i,j = 1, \ldots, m(\delta)\right\}. \tag{8}$$

The function $f_{\text{ol}}(\delta)$ counts the nodes that are overlapping at each value of δ, i.e. that belong to more than one block of the covering $K_X(\delta)$. From property (i) in Def. 7, it follows that $f_{\text{ol}}(0) = 0$ and $f_{\text{ol}}(\delta) = 0$ for all $\delta \geq \delta_{\max}$. Moreover, we use the overlapping function to define the notion of clusterability as follows.

Definition 12. *Given a network $N = (X, A_X)$, a cut metric c_X, and an associated nested covering K_X, we say that (X, c_X) is* clusterable *if there exists a δ such that $f_{ol}(\delta) = 0$, $K_X(\delta) \neq \{\{x\}, x \in X\}$, and $K_X(\delta) \neq \{X\}$.*

In other words, a network is clusterable if for some resolution δ we obtain a valid partition (i.e., there is no overlap) that is different from those partitions obtained for extreme resolutions ($\delta = 0$ where all nodes are clustered separately, and $\delta \geq \delta_{\max}$ where all nodes are clustered together). Even for non-clusterable networks, the overlapping function provides valuable information about the underlying grouping structure and may help to identify nodes that cannot be fully classified into one subset. Also, the overlapping function is useful in determining meaningful values of the resolution δ; see Sections VI and VII for additional details.

To sum up, in the same way that in hierarchical (non-overlapping) clustering, ultrametrics are used to define equivalence relations that determine dendrograms, in the proposed hierarchical overlapping clustering method, cut metrics are used to define tolerance relations that determine nested coverings; see Table I.

Remark 1. Recall that defining a partition on a node set is equivalent to defining an equivalence relation on the same set (cf. Defs. 1 and 2). This means that specifying any of them uniquely leads to the other. The same occurs for the definitions of dendrograms and ultrametrics; see Theorem 1. However, in the context of hierarchical overlapping clustering these bijections no longer hold. While it is true that given a cut metric, we can obtain a collection of tolerance relations that leads to a nested collection of coverings (see Theorem 2), we cannot, in general, uniquely determine a tolerance relation from a given covering [20, Theorem 2] or a cut metric from a given nested collection of coverings.

Remark 2. In a strict sense, the proof of Theorem 2 does not require c_X to be a cut metric. In fact, any nonnegative, symmetric function satisfying the identity property would suffice to create a relationship between nodes as in (7). Constructing such a relation is known as the Rips complex [26]. However, since we ultimately want to group nodes that are more similar to each other than to the rest, we are looking for functions that actually reflect the relationship between nodes encoded in the dissimilarity function. The cut metric obtained as a convex combination of ultrametrics achieves this. Moreover, due to its construction (cf. Figs. 2 and 3), cut metrics are intrinsically related to the formation of coverings and partitions. Another recent alternative to use in (7) is projections on A-spaces as done in [27].

IV. Hierarchical overlapping clustering algorithm

Given a network, our goal is to output a nested collection of coverings such that the grouping of nodes is done in accordance to the difference between them at varying levels of resolution. In virtue of Theorem 2, a cut metric can be used to induce a nested collection of coverings. Hence, the idea is to systematically construct a cut metric in such a way that it reflects the dissimilarity across nodes as given by the dissimilarity function A_X. The following proposition states a key relation between ultrametrics and cut metrics.

Proposition 1. *A convex combination of ultrametrics yields a cut metric.*

Proof: Given (X, d_X) where d_X is a metric, then d_X is in particular a tree metric – there exists a tree graph in which it is possible to embed the distance on its edges [18, p. 147] – if and only if it satisfies the four-point condition

$$d_X(x,x') + d_X(x'',x''') \leq \max\left\{d_X(x,x'') + d_X(x',x'''),\ d_X(x,x''') + d_X(x',x'')\right\} \tag{9}$$

for all $x, x', x'', x''' \in X$. Ultrametrics also satisfy the four-point condition so that they are particular cases of tree metrics [18, p. 311]. It can be shown that tree metrics, as well as convex combinations of tree metrics, are ℓ_1-embeddable [18, Proposition 11.1.4, Fact 11.1.15]. Finally, since a metric is ℓ_1-embeddable if and only if it is a cut metric [18, Proposition 4.4.2], a convex combination of ultrametrics yields a cut metric. ■

From Proposition 1 it immediately follows that ultrametrics are particular cases of cut metrics. Intuitively, this feature is inherited from the facts that equivalence relations are particular cases of tolerance relations and that partitions are particular cases of coverings; see Table I. Also, Proposition 1 plays a key role in the generation of cut metrics as it gives

Algorithm 1 Hierarchical overlapping clustering algorithm $\mathcal{O}$.

Input: J: no. of perturbations, N: network, $\mathcal{H}$: hierarchical clustering method, perturbation$(\cdot)$
Output: K_X: nested collection of coverings.

procedure $\mathcal{O}(J, N, \mathcal{H}, \text{perturbation})$
  for i$=1{:}J$ **do**
    $\tilde{N}$ = perturbation(N);
    $\tilde{u}_X^{(i)} = \mathcal{H}(\tilde{N})$;
  end for
  $c_X = \text{avg}(\tilde{u}_X^{(1)}, \ldots, \tilde{u}_X^{(J)})$;
  $K_X = \text{obtain_coverings}(c_X)$;
end procedure

a systematic way to obtain cut metrics from ultrametrics, which are readily available from the application of hierarchical (non-overlapping) clustering methods $\mathcal{H}$.

The next step is to obtain multiple ultrametrics related to the dissimilarity function A_X so that the resulting combination of these yields a cut metric that bears a relation with A_X. In order to achieve this, the concept of dithering [21] proves to be useful, where we intentionally apply random noise to the dissimilarity function A_X to obtain $\tilde{A}_X$. Denoting by $\tilde{N} = (X, \tilde{A}_X)$ the resulting perturbed network, and by $\tilde{u}_X = \mathcal{H}(\tilde{N})$ the ultrametric obtained by hierarchically clustering $\tilde{N}$, the following corollary of Proposition 1 results.

Corollary 1. *The function c_X obtained as*

$$c_X = \mathbb{E}\left[\tilde{u}_X\right] \tag{10}$$

is a cut metric, where the expectation is taken with respect to the probability distribution of the randomization introduced by dithering.

In practice, several realizations of random perturbations are implemented, generating a whole family of closely related networks. Each one of these networks yields a different ultrametric when a pre-specified hierarchical clustering method $\mathcal{H}$ is applied to them. When all these ultrametrics are averaged, a cut metric is obtained (cf. Proposition 1). Formally, given a network $N = (X, A_X)$ denote by $\tilde{N}_i = (X, \tilde{A}_X^{(i)})$ for $i = 1, \ldots, J$ the J networks resulting from independently perturbing the dissimilarity function J times. Moreover, denoting by $\tilde{u}_X^{(i)}$ the ultrametric output of hierarchically clustering $\tilde{N}_i$, we generate the cut metric c_X given by (cf. Proposition 1)

$$c_X(x, x') = \frac{1}{J} \sum_{i=1}^{J} \tilde{u}_X^{(i)}(x, x'). \tag{11}$$

Finally, we use c_X to obtain an associated nested collection of coverings K_X (cf. Theorem 2). The algorithm is summarized in Algorithm 1.

Remark 3. Proposition 1 also holds for non-negative linear combinations of ultrametrics, but only a convex combination guarantees that the distance scale described by the ultrametrics is preserved in the resulting cut metric. Therefore, if we want to interpret the scaling parameter δ in the same units as the ultrametric function, then we should proceed with a convex combination. This is useful for directly comparing how clusters are formed in hierarchical methods, both overlapping and non-overlapping.

Remark 4. Observe that cut metrics conform a convex cone and hence, any non-negative combination of cut metrics is still a cut metric. This is not true in the case of ultrametrics: a non-negative linear combination of ultrametrics is not an ultrametric [25, Section VII-B]. In fact, it is cut metrics that form the conic hull of ultrametrics.

V. Quasi-Cut Metrics

The output structure of the hierarchical overlapping clustering method $\mathcal{O}$ is a nested collection of coverings $\mathcal{K}$ (cf. Def 8). This is a symmetric structure in the sense that if node x shares a block in a covering with node x' then x' must share a block with x. This feature is inherited from the fact that cut metrics are symmetric functions. When considering asymmetric (directed) networks, the symmetry in the output necessarily entails a loss of information, which might be an undesirable feature of a grouping algorithm in some applications [28]–[30].

Our objective is to provide a framework for developing a hierarchical overlapping quasi-clustering of the node set X. A quasi-cluster is a structure that consists of both a grouping of the node set and a collection of edges indicating the influence between different groups of nodes [31]. In extending this notion, we start by defining quasi-coverings, which are the basic units of the proposed method, and we also determine the restrictions required for a collection of quasi-coverings to be nested.

Definition 13. *The pair $\breve{Q}_X = (Q_X, F_X)$ is a quasi-covering if $Q_X = \{C_1, \ldots, C_m\}$ is a covering of X and $F_X \subseteq Q_X \times Q_X$ is a set (possibly empty) of ordered pairs representing directed edges between the elements of Q_X.*

A quasi-covering is a directed graph whose node set is given by the blocks of a covering. The edges represent the asymmetric influence that some blocks can exercise over others. The conditions needed for a collection of quasi-coverings to be nested are described next.

Definition 14. *The collection $\breve{K}_X = \{\breve{K}_X(\delta), \delta \geq 0\}$ is a nested collection of quasi-coverings if $\breve{K}_X(\delta) = (K_X(\delta), F_X(\delta))$ is a quasi-covering for each resolution parameter δ, $K_X = \{K_X(\delta), \delta \geq 0\}$ is a nested collection of coverings (cf. Def. 7), and $F_X = \{F_X(\delta), \delta \geq 0\}$ is a collection of edge sets satisfying*

(i) $F_X(0) = \emptyset$ and there exists a $\delta_{\max}$ such that $F_X(\delta) = \emptyset$ for all $\delta \geq \delta_{\max}$;
(ii) For arbitrary resolutions $\delta < \delta'$ and covers $C_i, C_j \in K_X(\delta)$ such that $(C_i, C_j) \in F_X(\delta)$, there exist covers $C_i', C_j' \in K_X(\delta')$ such that $C_i \subseteq C_i'$, $C_j \subseteq C_j'$ and either $(C_i', C_j') \in F_X(\delta')$ or $|C_i \cap C_j| < |C_i' \cap C_j'|$.

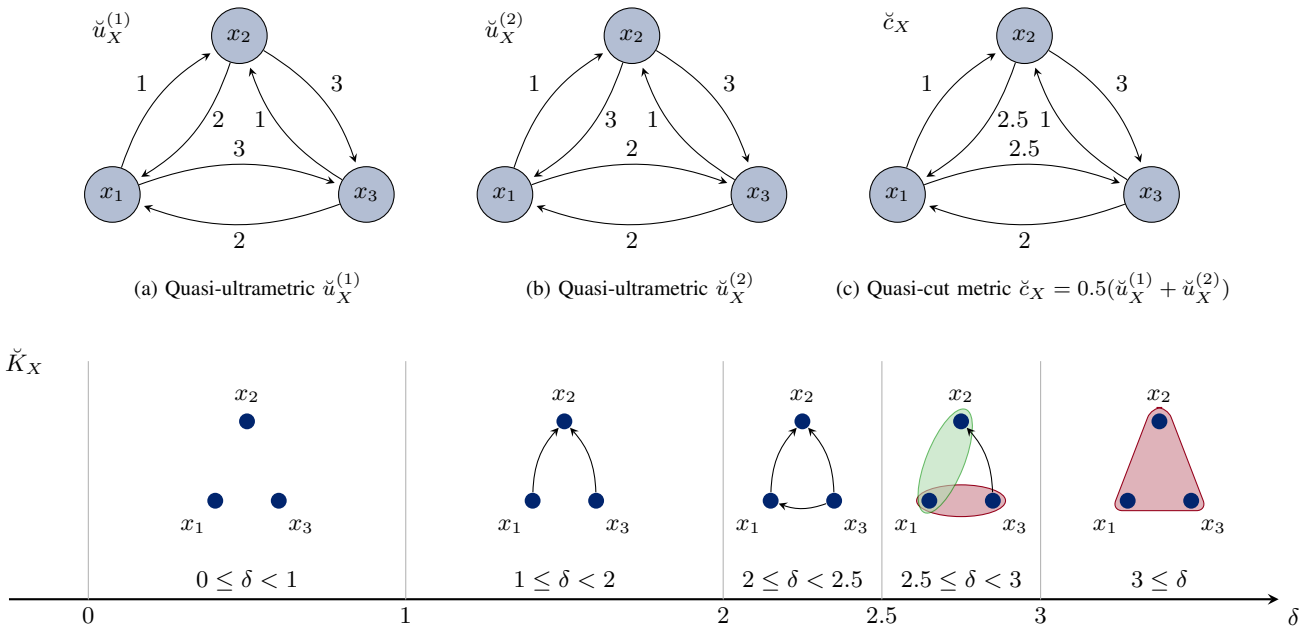

(a) Quasi-ultrametric $\breve{u}_X^{(1)}$ (b) Quasi-ultrametric $\breve{u}_X^{(2)}$ (c) Quasi-cut metric $\breve{c}_X = 0.5(\breve{u}_X^{(1)} + \breve{u}_X^{(2)})$

(d) Nested collection of quasi-coverings $\breve{K}_X$

Figure 4. Illustrative example of a hierarchical overlapping quasi-clustering method $\breve{\mathcal{O}}$. (a)-(b) Given quasi-ultrametrics. (c) Quasi-cut metric obtained by averaging the given quasi-ultrametrics. (d) Nested collection of quasi-coverings obtained from $\breve{c}_X$. For $\delta = 0$ all nodes belong to separate covers and the edge set is empty. At $\delta = 1$ nodes x_1 and x_3 start influencing x_2. At $\delta = 2$ node x_3 also exercises influence over x_1. The first two covers form at $\delta = 2.5$ and they present an overlap, since x_1 belongs to both covers. Additionally, node x_3 still exercises influence over x_2. Finally, for $\delta = 3$ all nodes belong to the same single cover and the edge set is empty.

Denote by $\breve{\mathcal{K}}$ the space of all nested collections of quasi-coverings.

A nested collection of quasi-coverings reflects the different levels of similarity present in the network. More specifically, as δ grows, covers would typically start exercising influence one over the other, then some nodes start to overlap while the remaining ones still exercise influence, until eventually the covers merge into one; see Fig. 4d for an example. Requirement (i) in Def. 14 enforces border conditions on the edge set F_X that are consistent with those required for K_X in Def. 7. More precisely, at resolution $\delta = 0$, each node must belong to its own cluster *and* there should be no influence relations among these. By contrast, for a large enough resolution $\delta_{\max}$ every node must belong to a single cluster, hence, no influence relation is possible. Requirement (ii) enforces nestedness in the influence structure F_X. Specifically, if an edge exists at resolution δ between blocks C_i and C_j then this edge must persist for larger resolutions δ' or the blocks must become more similar, encoded by having a larger intersection.

We consider a nested collection of quasi-coverings to be the desired output of hierarchical overlapping quasi-clustering methods. Formally, define such methods $\breve{\mathcal{O}}$ as a structure-preserving map from the space of networks to the space of nested collections of quasi-coverings, $\breve{\mathcal{O}} : \mathcal{N} \to \breve{\mathcal{K}}$. In analogy to the development of Sections III and IV, in which we used hierarchical (non-overlapping) clustering methods to obtain ultrametrics that are then combined to obtain cut metrics, in what follows we propose to use existing hierarchical (non-overlapping) quasi-clustering methods to obtain quasi-ultrametrics that can be combined into quasi-cut metrics (Def. 15). Quasi-ultrametrics are defined as non-negative functions $\breve{u}_X : X \times X \to \mathbb{R}_+$ such that $\breve{u}_X(x, x') = 0$ if and only if $x = x'$ and $\breve{u}_x(x, x') \le \max\{\breve{u}_X(x, x''), \breve{u}_X(x'', x')\}$ for all $x, x', x'' \in X$ but not necessarily symmetric $\breve{u}_X(x, x') \neq \breve{u}_X(x', x)$; see [31]. Denote by $\breve{\mathcal{U}}$ the set of all possible quasi-ultrametrics.

Definition 15. *Denote as $\breve{\mathcal{V}}$ the conic hull of the set of quasi-ultrametrics, $\breve{\mathcal{V}} = cone\{\breve{\mathcal{U}}\}$. Then, given a node set X, an element $\breve{c}_X \in \breve{\mathcal{V}}$ is a quasi-cut metric on the node set X.*

Notice that, by definition, a convex combination of quasi-ultrametrics outputs a quasi-cut metric. Moreover, nested collections of quasi-coverings can be constructed from quasi-cut metrics as shown in the following theorem.

Theorem 3. *Given a node set X and a quasi-cut metric $\breve{c}_X \in \breve{\mathcal{V}}$, we can obtain a nested collection of quasi-coverings $\breve{K}_X = \{\breve{K}_X(\delta), \delta \ge 0\}$, with $\breve{K}_X(\delta) = (K_X(\delta), F_X(\delta))$ a quasi-covering, as follows. The covering $K_X(\delta)$ is obtained from the tolerance relation $\leftrightarrow_\delta$ determined by*

$$\max\{\breve{c}_X(x, x'), \breve{c}_X(x', x)\} \le \delta \Rightarrow x \leftrightarrow_\delta x'. \tag{12}$$

The edge set $F_X(\delta)$ is constructed such that for $C_i, C_j \in$

$K_X(\delta)$, we have that $(C_i, C_j) \in F_X(\delta)$ if

$$\min_{\substack{x \in C_i \backslash C_i \cap C_j \\ x' \in C_j \backslash C_i \cap C_j}} \breve{c}_X(x, x') \leq \delta. \tag{13}$$

Proof. Notice that $\max\{\breve{c}_X(x, x'), \breve{c}_X(x', x)\}$ is a nonnegative, symmetric function that satisfies the identity property. Hence, $\leftrightarrow_\delta$ in (12) is a valid tolerance relation and, from the combination of Theorem 2 and Remark 2, it follows that $K_X = \{K_X(\delta), \delta \geq 0\}$ is a nested collection of coverings as required by Def. 14.

We now show that the collection of edge sets F_X satisfies conditions (i) and (ii) in Def. 14. It follows from the definition of quasi-cut metrics that $\breve{c}_X(x, x') = 0$ if and only if $x = x'$ so that $F_X(0) = \emptyset$. Also, since K_X is a nested collection of coverings, it must be that for some $\delta_{\max}$ we have $K_X(\delta) = \{X\}$ for all $\delta \geq \delta_{\max}$. Consequently, having only one block, the edge set must be empty, i.e., $F_X(\delta) = \emptyset$ for all $\delta > \delta_{\max}$.

To prove that F_X satisfies condition (ii), let $C_i, C_j \in K_X(\delta)$ and $(C_i, C_j) \in F_X(\delta)$ at a given δ. Then, there is (at least) a pair of nodes $x_i \in C_i \backslash C_i \cap C_j$ and $x_j \in C_j \backslash C_i \cap C_j$ that satisfy (13). Focus now on an arbitrary resolution $\delta' > \delta$. The fact that $K_X = \{K_X(\delta), \delta \geq 0\}$ is a nested collection of coverings implies the existence of blocks $C_i', C_j' \in K_X(\delta')$ such that $C_i \subseteq C_i'$ and $C_j \subseteq C_j'$. This immediately implies that $x_i \in C_i'$ and $x_j \in C_j'$. Consider then the following two alternatives: (1) if neither x_i nor x_j belong to $C_i' \cap C_j'$, then (13) forces $(C_i', C_j') \in F_X(\delta')$; and (2) if either x_i or x_j belong to $C_i' \cap C_j'$, then we have that $|C_i' \cap C_j'| > |C_i \cap C_j|$. Given that for any δ' either (1) or (2) must be true, requirement (ii) in Def. 14 is satisfied, completing the proof. □

In virtue of Theorem 3 we can readily obtain a nested collection of quasi-coverings from a quasi-cut metric. The quasi-cut metric can thus be constructed by performing a conic combination of quasi-ultrametrics. The algorithm proposed is analogous to the one developed in Section IV. Given a network (X, A_X), first dither J times the dissimilarity function A_X. Then, obtain a quasi-ultrametric $\breve{u}_X^{(i)}$ by applying a hierarchical (non-overlapping) quasi-clustering method – see, e.g., [31] – to each of the dithered networks $\tilde{N}_i = (X, \tilde{A}_X^{(i)})$ for $i = 1, \ldots, J$. Finally, generate the quasi-cut metric $\breve{c}_X$ by averaging the J quasi-ultrametrics.

In Fig. 4 we show an illustrative example of a hierarchical overlapping quasi-clustering method $\breve{\mathcal{O}}$. We are given two quasi-ultrametrics $\breve{u}_X^{(1)}$ and $\breve{u}_X^{(2)}$ – representing those obtained by dithering the same underlying network – depicted in Figs. 4a and 4b, respectively. We then obtain a quasi-cut metric by averaging these two quasi-ultrametrics $\breve{c}_X = 0.5(\breve{u}_X^{(1)} + \breve{u}_X^{(2)})$; see Fig. 4c. Observe that $\breve{c}_X$ is not a quasi-ultrametric since $\breve{c}_X(x_2, x_3) = 3$ which is larger than $\max\{\breve{c}_X(x_2, x_1), \breve{c}_X(x_1, x_3)\} = 2.5$. In Fig. 4d we portray the nested collection of quasi-coverings resulting from the quasi-cut metric $\breve{c}_X$ as explained in Theorem 3. It is observed that for $\delta = 0$ all nodes belong to separate covers and the edge set is empty; and that for $\delta = 3$ all nodes belong to the same single cover and the edge set is also empty [cf. (i) in Def. 14]. We can also observe the progression of influence as δ grows larger. Nodes start influencing each other ($\delta \geq 1$), until covers are created ($\delta = 2.5$) based on nodes that were already exercising influence on each other. Then, some nodes overlap (x_1), while non-overlapping nodes still exercise influence (x_3 onto x_2). Finally, all covers and edges collapse into a single cover containing all nodes ($\delta \geq 3$).

Remark 5. Quasi-coverings in Def. 13 act as generalizations of quasi-partitions introduced in [31]. Likewise, a nested collection of quasi-coverings is the generalization of a quasi-dendrogram. These natural extensions are obtained from the constructions in [31] by dropping the non-intersecting requirement.

VI. Synthetic Experiments

The first three illustrations of Algorithm 1 are conducted on synthetic networks $N = (X, A_X)$ where X is a set of points embedded in $\mathbb{R}^2$ and A_X is the Euclidean distance between them. As a hierarchical clustering method $\mathcal{H}$ we apply single linkage [13] and average the clustering output of $J = 100$ different noisy realizations of N. The noisy $\tilde{N}$ are obtained by perturbing the positions of the nodes in $\mathbb{R}^2$ with zero-mean gaussian noise with standard deviation σ obtained as 10^{-1} of the smallest distance between nodes unless otherwise stated.

A. Two clouds

Consider the simple network portrayed in Fig. 5a, where there are 121 nodes on each cloud. The average distance between adjacent nodes within each of the clouds is $d_1 = 1$ whereas the distance between the centers of both clouds is $d_2 = 22$. The overlapping function is found in Fig. 5e. As expected, this network is clusterable (see Definition 12), i.e., for $\delta = 1.11$ a non-trivial (neither all nodes together nor all separated) partition $K_X(\delta) = \{C_1, C_2\}$ is obtained; see Fig. 5a. Observe that the value of δ that generates this partition that identifies both clouds is in the order of d_1.

B. Multiple resolutions

Consider now the network depicted in Fig. 5b which consists of five clouds of 81 points each. Within each of the four clouds on the left, the average distance between adjacent nodes is $d_1 = 1$ whereas the average distance between neighboring nodes in different clouds is $d_2 = 2$. The fifth cloud, on the right, has an average distance between adjacent nodes within the cloud of d_2 and the distance between the center of the right cloud and that of the other four clouds is $d_3 = 22$. The resulting overlapping function is found in Fig. 5f.

The network is found to be clusterable. First, for $\delta = 1.11$ we obtain the covering $K_X(\delta) = \{C_1, C_2, \ldots, C_{85}\}$, as illustrated in Fig. 5b. In this case, each subset $C_1, \ldots, C_4$ contains one of the clouds on the left and each of the subsets $C_5, \ldots, C_{85}$ contains a single point of the fifth cloud. Intuitively, since $\delta = 1.11$ is close to d_1 but smaller than

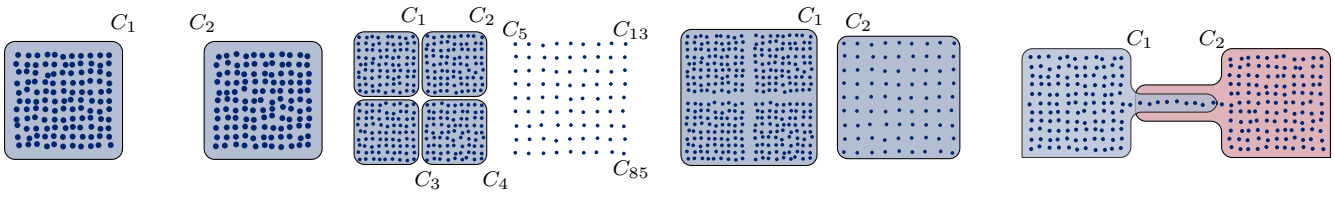

(a) Two clouds (covers for $\delta = 1.11$)

(b) Multiple resolutions (covers for $\delta = 1.11$)

(c) Multiple resolutions (covers for $\delta = 2.07$)

(d) Dumbbell (covers for $\delta = 2.17$)

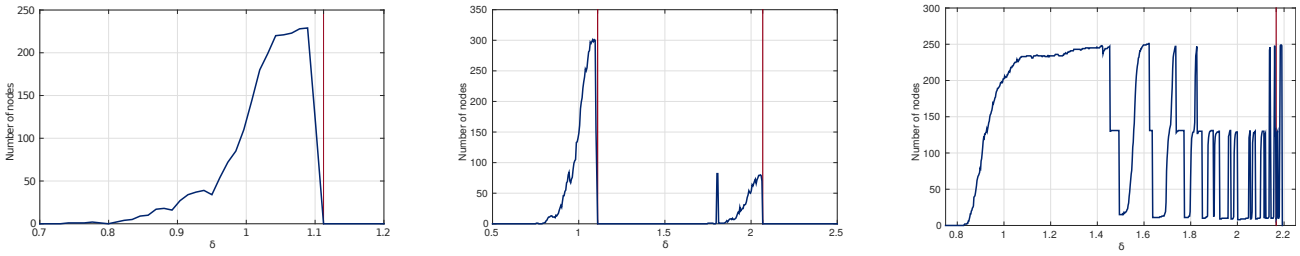

(e) Two clouds: overlapping function

(f) Multiple resolutions: overlapping function

(g) Dumbbell: overlapping function

Figure 5. Synthetic experiments. (a)-(d) Location of nodes for the three synthetic networks considered: (a) two clouds, (b)-(c) multiple resolutions, and (d) dumbbell. Also shown are the coverings obtained at different values of δ. (e)-(g) Overlapping functions for each one of the mentioned networks. Shown in red are the values of δ for which either a zero – (e) and (f) – or a minimum of local minima – (g) – is attained.

d_2, only nodes that are at distances around or below d_1 are grouped together.

Additionally, the network is also clusterable for $\delta = 2.07$ with the resulting covering $K_X(\delta) = \{C_1, C_2\}$ portrayed in Fig. 5c. In this case, there are only two subsets: C_1 contains the four clouds on the left, and C_2 contains the right cloud. This is reasonable since $\delta = 2.07$ is slightly greater than both the inter-cloud distance on the left and the inter-point distance on the right but smaller than the minimum distance between points in the left and right clouds.

C. A solution to single linkage's chaining effect

Consider the network formed by the set of nodes in Fig. 5d where the average distance between adjacent nodes is $d = 1$. For this simulation, the value of σ used is equal to the minimum distance between nodes. Unlike the previous examples considered, the expected clustering output is not unequivocal since it is not clear how to group the nodes along the line connecting both point clouds. More specifically, it is not clear whether these points should constitute a cluster in itself or if this putative group should have overlap with the ones corresponding to the point clouds at its extremes. As a matter of fact, if one applies (non-overlapping) single linkage clustering to the (unperturbed) network of interest, the output dendrogram only consists of the two trivial partitions: every node in a different cluster or all nodes clustered together. This phenomenon is known as chaining effect [9] and is an undesirable concomitant of the definition of single linkage. Nevertheless, via the introduction of dithering, non-trivial coverings of the network of interest can be recovered, thus, overcoming the aforementioned chaining effect.

The overlapping function is presented in Fig. 5g and, in accordance with the intuition explained, the network is found not to be clusterable, i.e., no zeros are attained for intermediate resolutions. In this case, we focus on the resolution δ leading to the minimum number of overlapping nodes, i.e., we look for a covering associated with the minimum of all local minima of the overlapping function. For $\delta = 2.17$, the resulting covering is depicted in Fig. 5d where there are only two subsets C_1 and C_2, one containing all the left cloud as well as the connecting line (except for the rightmost point in the line), while the other subset contains all the right cloud and the connecting line (except for the leftmost point). By admitting overlap between clusters and using the overlapping function to guide our analysis, useful insight on the structure of the network was recovered.

VII. Applications

We consider two real-world classification problems: determining handwritten digits from the MNIST image database (Section VII-A), and identifying the authors and co-authors of famous plays (Section VII-B).

A. Handwritten digit classification

The proposed hierarchical overlapping clustering algorithm is implemented to classify digits of the MNIST database [22]. The database consists of black and white images $\{\mathbf{G}_i\}_{i\in\mathcal{I}}$ of size 28×28 pixels.

In order to construct a network from the images, we first reshape each image $\mathbf{G}_i$ into a column vector $\mathbf{y}_i$ of size 784 and then apply a principal component analysis (PCA) transform [32] to each vector. For computing the PCA, $5,000$ training samples of each digit are used to estimate the mean $\hat{\boldsymbol{\mu}}$ and the covariance matrix $\hat{\mathbf{C}}$. The resulting

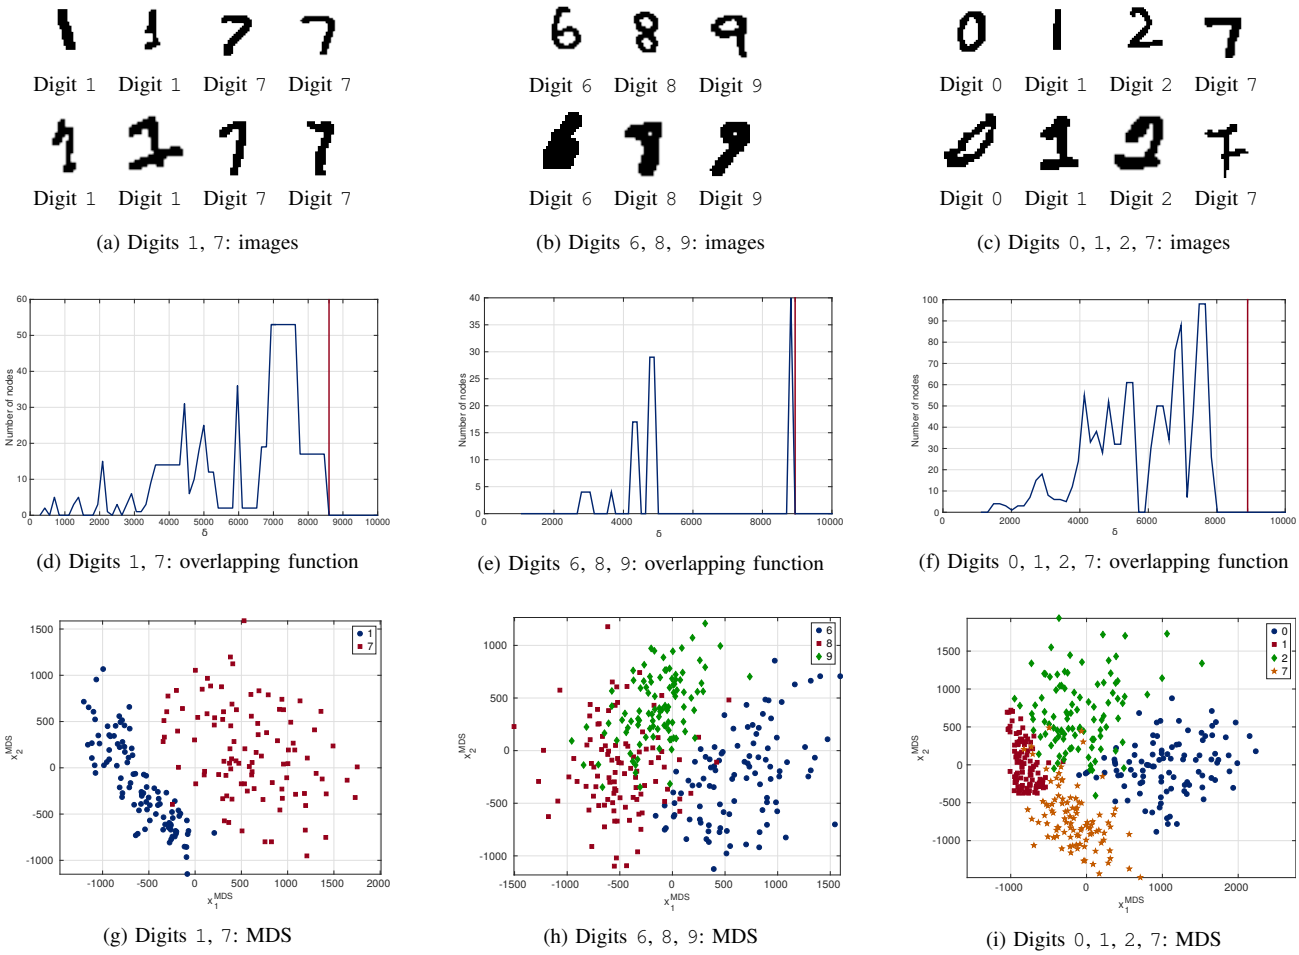

(a) Digits 1, 7: images (b) Digits 6, 8, 9: images (c) Digits 0, 1, 2, 7: images

(d) Digits 1, 7: overlapping function (e) Digits 6, 8, 9: overlapping function (f) Digits 0, 1, 2, 7: overlapping function

(g) Digits 1, 7: MDS (h) Digits 6, 8, 9: MDS (i) Digits 0, 1, 2, 7: MDS

Figure 6. Classification of handwritten digits. (a)-(c) Samples of images of the digits for each one of the three classification problems: (a) digits 1 and 7, (b) digits 6, 8, and 9, (c) digits 0, 1, 2, and 7. The first row shows samples of numbers that are typically classified correctly by the proposed method whereas the second row presents numbers that are hard to classify. (d)-(f) Overlapping functions for each of the classification problems considered. (g)-(i) Multidimensional Scaling (MDS) representation in 2 dimensions of the points in the 20-PCA domain.

transformed vectors are denoted $\mathbf{y}_i^{\text{PCA}}$ for $i \in \mathcal{I}$. A total of n images are considered for classification, and only the first 20 components of the associated transformed vectors $\mathbf{y}_i^{\text{PCA}}$ are used while the rest are discarded. The reduced vectors $\mathbf{x}_i \in \mathbb{R}^{20}$ for $i = 1, \ldots, n$ constitute the node set $X = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ of our network. Each element of the (symmetric) dissimilarity function is computed as the Euclidean distance between vectors in the 20-PCA space, i.e., $A_X(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|_2$.

In the application of Algorithm 1, $J = 100$ noise realizations are carried out for the dithering step in all simulations. The standard deviation of the noise $\sigma = k \min_{\mathbf{x}_i \neq \mathbf{x}_j} A_X(\mathbf{x}_i, \mathbf{x}_j)$ is computed as a factor k of the minimum distance between nodes. The hierarchical (non-overlapping) clustering method $\mathcal{H}$ used to obtain the ultrametrics is Ward's clustering method [11].

1) Two-digit classification. Digits 1 and 7: First, the performance of the algorithm is analyzed for classification of just two digits: 1 and 7. The total number of images considered is $n = 200$ counting 100 images of each digit, selected at random from the test set of the MNIST database. Images of the digits 1 and 7 from this database are illustrated in Fig. 6a, where the top row shows typical images that are correctly classified when using the proposed method, and the bottom row shows numbers that are often confused.

The overlapping function obtained is displayed in Fig. 6d for $k = 0.01$ in the computation of σ. For $\delta = 8600$, the covering $K_X(\delta) = \{C_1, C_2\}$ with $C_1 = \{1(\times 100), 7\}$ and $C_2 = \{7(\times 99)\}$ is obtained. In this case, the proposed algorithm makes one classification mistake, hence, the error rate is 0.5% (1 misclassified image out of 200). To help illustrate this example, a non-metric Multidimensional Scaling (MDS) [33] is performed on the vectors in $\mathbb{R}^{20}$ to be able to depict them on $\mathbb{R}^2$. This representation is shown in Fig. 6g. Notice that points representing images of different digits are clearly separated, except for one instance of 7 that is projected on top of the points corresponding to digit 1.

2) Three-digit classification. Digits 6, 8 and 9: As done before, 100 images are chosen at random from the test dataset for each digit, resulting in $n = 300$ images. As can be observed from Fig. 6b that contains images of

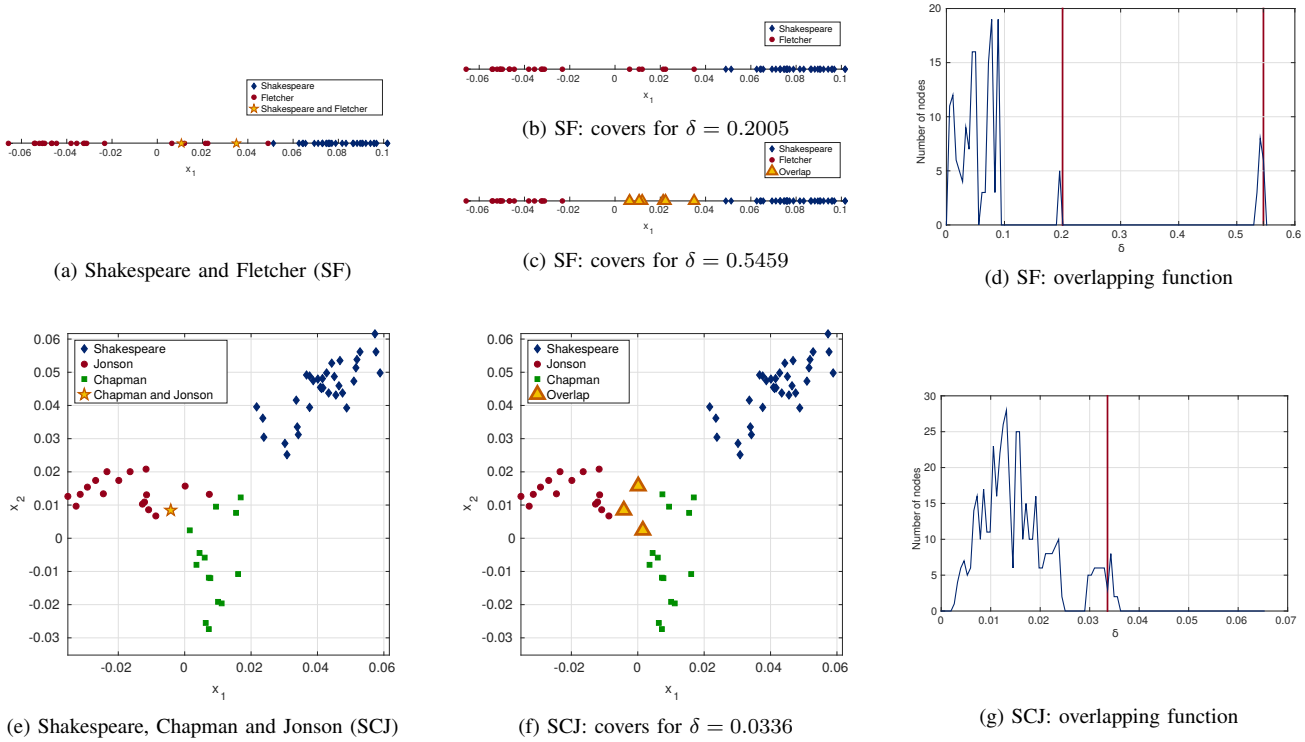

(a) Shakespeare and Fletcher (SF)

(b) SF: covers for $\delta = 0.2005$

(c) SF: covers for $\delta = 0.5459$

(d) SF: overlapping function

(e) Shakespeare, Chapman and Jonson (SCJ)

(f) SCJ: covers for $\delta = 0.0336$

(g) SCJ: overlapping function

Figure 7. Authorship attribution. (a)-(d) Classifying plays from Shakespeare and Fletcher. (a) Location of plays with correct labeling. (b) Covers obtained at resolution $\delta = 0.2005$ showing no overlap. (c) Covers for $\delta = 0.5459$ present overlap between Shakespeare and Fletcher. (d) Overlapping function showing the two resolutions of interest. (e)-(g) Classifying plays from Shakespeare, Chapman and Jonson. (e) Location of plays with correct labeling. (f) Covers obtained for $\delta = 0.0336$ showing overlap between Chapman and Jonson. (g) Overlapping function highlighting the local minimum of interest.

the numbers being classified, digits 8 and 9 are much harder to distinguish than digit 6. The overlapping function is portrayed in Fig. 6e for $k = 0.001$. The covering $K_X(\delta) = \{C_1, C_2, C_3\}$ for $\delta = 8941$ yields an error rate of 2.67% since $C_1 = \{6(\times 100)\}$, $C_2 = \{8(\times 95), 9(\times 3)\}$ and $C_3 = \{8(\times 5), 9(\times 97)\}$. The MDS representation is depicted in Fig. 6h, confirming that the major concern is separating 8 and 9 since these two digits occupy approximately half of the plane and 6 occupies the other half.

3) Four-digit classification. Digits 0, 1, 2 and 7: We again consider 100 images per digit totalizing $n = 400$ images randomly selected from the test dataset. Examples of these digits are in Fig. 6c and the plot of the overlapping function can be found in Fig. 6f for $k = 0.005$. For $\delta = 8901$ the resulting covering is $K_X(\delta) = \{C_1, C_2, C_3, C_4\}$ with $C_1 = \{0(\times 100), 2(\times 3)\}$, $C_2 = \{1(\times 99), 7(\times 2)\}$, $C_3 = \{1, 2(\times 86), 7(\times 3)\}$, and $C_4 = \{2(\times 11), 7(\times 95)\}$, resulting in an error rate of 5%. The MDS representation is shown in Fig. 6i where we see that 2 is as sparsely located as 0 but presents larger overlap with the clouds corresponding to 1 and 7. This constitutes the main source of error as can be seen by inspecting blocks C_3 and C_4.

B. Authorship Attribution

We address the problem of authorship attribution [23], where our objective is to attribute a given play to its rightful author. The authors considered in this experiment are William Shakespeare, John Fletcher, Ben Jonson, and George Chapman. We consider 33 plays written by Shakespeare, 21 by Fletcher, 17 by Jonson, 14 by Chapman, as well as 2 plays co-authored by Shakespeare and Fletcher and 1 co-authored by Jonson and Chapman. Following the procedure described in [23] based on word adjacency networks, for each play we are able to obtain a dissimilarity measure to the four authors of interest. Denote by y_{ij} the dissimilarity of play i to author j. The index j represents the initial of each author, $j \in \{S, F, C, J\}$. To be more specific, if y_{ij} is small, then play i follows a function-word structure that resembles the one typically used by author j; see [23] for details. In what follows, the node set X is comprised of the subset of plays corresponding to the authors under study. Our goal is to cluster together plays written by the same author, as well as identifying co-authored plays in the overlap between author-specific clusters.

1) Shakespeare and Fletcher: Denote by $\mathcal{S}$ the set of 33 plays written solely by Shakespeare, by $\mathcal{F}$ the set of 21 plays solely authored by Fletcher and by $\mathcal{SF}$ the set containing the two co-authored plays. Define by x_i the difference of dissimilarities of play i to each author, $x_i = y_{iF} - y_{iS}$, for all $i \in \mathcal{S} \cup \mathcal{F} \cup \mathcal{SF}$. The node set is then comprised of these 56 points on $\mathbb{R}$ representing the plays under study, i.e., $X = \{x_i\}_{i \in \mathcal{S} \cup \mathcal{F} \cup \mathcal{SF}}$. The location of these points is depicted in Fig. 7a. The dissimilarity function considered is the Euclidean distance $A_X(x_k, x_\ell) = |x_k - x_\ell|$ for every pair of plays k and l.

In the application of Algorithm 1, we use Ward's linkage

[11] as the method $\mathcal{H}$, we set $J = 100$ dithering steps, and we fix $k = 4$ in the determination of the noise amplitude; see Section VII-A. The resulting overlapping function is displayed in Fig. 7d and two coverings of interest achieved for resolutions $\delta = 0.2005$ and $\delta = 0.5459$ are presented in Figs. 7b and 7c, respectively.

In Fig. 7b it is observed that the data does not have any overlap and that there are two resulting clusters $C_1 = \{$Shakespeare (x 33), Fletcher$\}$ and $C_2 = \{$Fletcher (x 20), Shakespeare and Fletcher (x 2)$\}$. Thus, as a non-overlapping clustering method, it is seen that the proposed algorithm makes only one mistake but fails to identify the co-authored plays, assigning them to just one of the authors. By contrast, in Fig. 7c the coverings $C_1 = \{$Shakespeare (x 33), Fletcher (x 5), Shakespeare and Fletcher (x 2)$\}$ and $C_2 = \{$Fletcher (x 20), Shakespeare and Fletcher (x 2)$\}$ *do* overlap. More specifically, one error is committed and 6 plays belong to the overlap, among which the two co-authored plays are contained.

2) Shakespeare, Chapman, and Jonson: For the last application example, plays by Shakespeare, Chapman and Jonson are classified, including one play co-authored by the last two authors. Denote by $\mathcal{S}$, $\mathcal{C}$, $\mathcal{J}$, and $\mathcal{CJ}$, the sets containing the plays solely authored by Shakespeare, Chapman, and Jonson, and the co-authored play, respectively. Mimicking the construction in the previous experiment, for a given play i let $x_{i1} = y_{iC} - y_{iS}$ and $x_{i2} = y_{iJ} - y_{iS}$ be the differences in similarities between Chapman and Shakespeare, and Jonson and Shakespeare, respectively. Define the vector $\mathbf{x}_i = [x_{i1}, x_{i2}]^T \in \mathbb{R}^2$ for each play i to obtain the node set $X = \{\mathbf{x}_i\}_{i \in \mathcal{S} \cup \mathcal{C} \cup \mathcal{J} \cup \mathcal{CJ}}$ containing the 65 plays from all three authors. The position of these nodes in $\mathbb{R}^2$ is illustrated in Fig. 7e. In this case, the dissimilarity function is given by the Euclidean distance between the points, i.e., $A_X(\mathbf{x}_k, \mathbf{x}_\ell) = \|\mathbf{x}_k - \mathbf{x}_\ell\|_2$, for plays k and l.

In running Algorithm 1 we select UPGMA [10] as the hierarchical method $\mathcal{H}$, we set $J = 100$ dithering steps and the noise amplitude is chosen for $k = 1$. The resulting overlapping function is depicted in Fig. 7g and the covering for $\delta = 0.0336$ is displayed in Fig. 7f. This covering is comprised of the blocks $C_1 = \{$Shakespeare (x 33)$\}$, $C_2 = \{$Jonson (x 16), Chapman and Jonson$\}$, $C_3 = \{$Jonson (x 2), Chapman (x 14)$\}$ and $C_4 = \{$Jonson, Chapman, Chapman and Jonson$\}$. Note that block C_1 contains all of Shakespeare's plays and nothing else, thus making this classification perfect. Block C_2 contains all Jonson plays but one and also the co-authored play, and C_3 contains all Chapman plays and 2 Jonson plays, one of which overlaps with C_2. If these were all the outputs, then the algorithm would be making one full mistake (Jonson play classified as Chapman's), the co-authored play would not be recognized (as it is classified as Jonson's) and attention would be drawn to a Jonson play that is being classified as both Chapman and Jonson. However, the algorithm outputs a fourth covering C_4 that contains one play by Jonson (*"Sejanus"*), one play by Chapman (*"May Day"*) and the co-authored play (*"Eastward Ho"*). Thus, by outputting this fourth cluster, now there are three overlapping plays: one by each author and the co-authored one. Therefore, the algorithm successfully recognizes the co-authored play and draws attention to two other plays that are hard to classify.

VIII. Conclusions

The proposed hierarchical overlapping clustering method obtains a collection of coverings of the data in such a way that a nested structure exists among the coverings and nodes can simultaneously belong to more than one cluster. The essence of the method is drawn from the connection of cut metrics, tolerance relations, and nested collections of coverings, that are respective generalizations of ultrametrics, equivalence relations, and dendrograms, typical of hierarchical (non-overlapping) clustering methods. A systematic method for obtaining cut metrics by dithering the dissimilarity function of a network is also proposed. The overlapping function is introduced as a tool for gaining insight on the structure of the data. It is also helpful in identifying meaningful resolutions within the nested collection of coverings. Additionally, a hierarchical overlapping *quasi*-clustering method was proposed. The objective of such method is to design a grouping algorithm that accommodates for asymmetries in the data, inherent to directed networks. We devised an algorithm for constructing quasi-cut metrics and showed how to obtain the output of the quasi-clustering method from these.

The hierarchical overlapping clustering method was applied to three synthetic networks to establish its performance in controlled and intuitive scenarios, as well as to illustrate how overlap can solve single linkage's chaining effect. Moreover, the method was applied to classification of handwritten digits obtained from the MNIST database showing satisfactory results in clusterable datasets. Finally, the proposed algorithm was used to classify plays by author, and it was shown that it succeeds in recognizing co-authored plays.

References

[1] A. K. Jain and R. C. Dubes, *Algorithms for Clustering Data*, ser. Prentice Hall Advanced Reference Series. Englewood Cliffs, NJ: Prentice Hall, 1988.

[2] V. Estivill-Castro, "Why so many clustering algorithms: A position paper," *SIGKDD Explor. Newsl.*, vol. 4, no. 1, pp. 65–75, June 2002.

[3] R. B. Cattell, "The description of personality: Basic traits resolved into clusters," *J. Abnormal and Social Psychology*, vol. 38, no. 4, pp. 476–506, Oct. 1943.

[4] C. Lu, X. Hu, and J. Park, "Exploiting the social tagging network for web clustering," *IEEE Trans. Syst., Man, Cybern. A*, vol. 41, no. 5, pp. 840–852, Sep. 2011.

[5] M. Paulus and L. Kristoufek, "Worldwide clustering of the corruption perception," *Physica A: Statist. Mech. Appl.*, vol. 428, pp. 351–358, 2015.

[6] A. Ozdemir, M. Bolaños, E. Bernat, and A. Selin, "Hierarchical spectral consensus clustering for group analysis of functional brain networks," *IEEE Trans. Biomed. Eng.*, vol. 62, no. 9, pp. 2158–2169, Sep. 2015.

[7] L. A. Dalton, M. E. Benalcazar, M. Brun, and E. R. Dougherty, "Analytic representation of Bayes labeling and Bayes clustering operators for random labeld point processes," *IEEE Trans. Signal Process.*, vol. 63, no. 6, pp. 1605–1620, March 2015.

[8] X. Zao and A. H. Sayed, "Distributed clustering and learning over networks," *IEEE Trans. Signal Process.*, vol. 63, no. 13, pp. 3285–3300, July 2015.

[9] G. N. Lance and W. T. Williams, "A general theory of classificatory sorting strategies: 1. Hierarchical systems," *Comput. J.*, vol. 9, no. 4, pp. 373–380, Apr. 1967.

[10] R. R. Sokal and C. D. Michener, "A statistical method for evaluating systematic relationships," *University Kansas Sci. Bulletin*, vol. 38 (II), no. 22, pp. 1409–1438, March 1958.

[11] J. H. Ward Jr., "Hierarchical grouping to optimize an objective function," *J. Amer. Statist. Assoc.*, vol. 58, no. 301, pp. 236–244, March 1963.

[12] D. Defays, "An efficient algorithm for a complete link method," *Comput. J.*, vol. 20, no. 4, pp. 364–366, Apr. 1977.

[13] G. E. Carlsson and F. Mémoli, "Characterization, stability and convergence of hierarchical clustering methods," *J. Mach. Learning Res.*, vol. 11, pp. 1425–1470, Apr. 2010.

[14] M. A. Youssef, A. Youssef, and M. F. Younis, "Overlapping multihop clustering for wireless sensor networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 20, no. 12, pp. 1844–1856, Dec. 2009.

[15] P. C. H. Ma and K. C. C. Chan, "A novel approach for discovering overlapping clusters in gene expression data," *IEEE Trans. Biomed. Eng.*, vol. 56, no. 7, pp. 1803–1809, July 2009.

[16] J. C. Bezdek, *Pattern Recognition with Fuzzy Objective Function Algorithms*, ser. Advanced Appl. Pattern Recognition. New York, NY: Plenum Press, 1981.

[17] G. Cleuziou, "An extended version of the k-means method for overlapping clustering," in *Int. Conf. Pattern Recognition 2008*. Tampa, FL: Int. Assoc. Pattern Recognition, 8-11 Dec. 2008.

[18] M. Deza and M. Laurent, *Geometry of Cuts and Metrics*. Paris, France: LIENS - Ecole Normale Supérieure, Nov. 1996.

[19] J. Culbertson, D. P. Guralnik, and P. F. Stiller, "Injective metrizability and the duality theory of cubings," *arXiv:1502.00126v1 [math.MG]*, 31 Jan. 2015. [Online]. Available: http://arxiv.org/abs/1502.00126

[20] W. Bartol, J. Miró, J. Pióro, and F. Rosselló, "On the coverings by tolerance classes," *Inform. Sci.*, vol. 166, no. 1-4, pp. 193–211, Oct. 2004.

[21] L. Schuchman, "Dither signals and their effect on quantization noise," *IEEE Trans. Commun. Technol.*, vol. 12, no. 4, pp. 162–165, Dec. 1964.

[22] Y. Le Cun, C. Cortes, and C. J. C. Burges, "The MNIST database of handwritten digits," Website, 18 Aug. 2015. [Online]. Available: http://yann.lecun.com/exdb/mnist/

[23] S. Segarra, M. Eisen, and A. Ribeiro, "Authorship attribution through function word adjacency networks," *IEEE Trans. Signal Process.*, vol. 63, no. 20, pp. 5464–5478, Oct. 2015.

[24] P. R. Halmos, *Naive Set Theory*, ser. University Series Undergraudate Math. New York, NY: Van Nostrand Reinhold Company, 1960.

[25] G. E. Carlsson, F. Mémoli, A. Ribeiro, and S. Segarra, "Axiomatic construction of hierarchical clustering in asymmetric networks," *arXiv:1301.7724v2 [cs.LG]*, 2 Sep. 2014. [Online]. Available: http://arxiv.org/abs/1301.7724

[26] J.-C. Hausmann, "On the Vietoris-Rips complexes and a cohomology theory for metric spaces," in *Prospects in Topology*, ser. Ann. Math. Stud. Princeton, NJ: Princeton University Press, 1995, vol. 138, pp. 175–188.

[27] J. Culbertson, D. P. Guralnik, and P. F. Stiller, "Functorial hierarchical clustering with overlaps," *arXiv:1609.02513v1 [cs.LG]*, 8 September 2016. [Online]. Available: http://arxiv.org/abs/1609.02513

[28] P. B. Slater, "A partial hierarchical regionalization of 3140 US counties on the basis of 1965-1970 intercounty migration," *Environment and Planning A*, vol. 16, no. 4, pp. 545–550, Apr. 1984.

[29] W. Pentney and M. Meila, "Spectral clustering of biological sequence data," in *20th Nat. Conf. Artificial Intell.* Pittsburgh, PA: Assoc. Advancement Artificial Intell., 9-13 July 2005.

[30] Y. Zhao and G. Karypis, "Hierarchical clustering algorithms for document datasets," *Data Mining and Knowledge Discovery*, vol. 10, no. 2, pp. 141–168, March 2005.

[31] G. Carlsson, F. Mémoli, A. Ribeiro, and S. Segarra, "Hierarchical quasi-clustering methods for asymmetric networks," in *31st Int. Conf. Mach. Learning*. Beijing, China: J. Mach. Learning Res., 21-26 June 2014.

[32] J. E. Jackson, *A User's Guide to Principal Components*. New York, NY: John Wiley and Sons, 1991.

[33] J. B. Kruskal and M. Wish, *Multidimensional Scaling*, ser. Quantitative Appl. Social Sci. Iowa City, IA: Sage Publications, Inc., 1978.